

WHAT IS CLAIMED IS:

- 1 1. A storage system comprising:
2 a first interface configured to receive block-level I/O requests;
3 a second interface configured for file-level I/O;
4 a third interface configured for communication with another storage
5 controller; and
6 a data communication path suitable for data communication with physical
7 storage in order to exchange data with the physical storage in response to the block-level I/O
8 request and the file-level I/O requests,
9 wherein the block-level I/O requests are serviced by exchanging data with a
10 first storage area of the physical storage and the file-level I/O requests are serviced by
11 exchanging data with a second storage area of the physical storage,
12 wherein the first storage area is accessed as a first logical volume and the
13 second storage area is accessed as a second logical volume,
14 wherein the first logical volume and the second logical volume define a
15 consistency group,
16 wherein data contained in the first logical volume and in the second logical
17 volume can be communicated to the other storage controller in order to replicate the data to
18 secondary storage accessible by the other storage controller,
19 wherein time consistency of write operations made to the first logical volume
20 and to the second logical volume is maintained at the secondary storage.

- 1 2. The storage system as recited in claim 1 wherein the physical storage
2 comprises a plurality of physical storage devices.

- 1 3. A storage system comprising:
2 a storage controller;
3 a plurality of physical disk drives connected to the storage controller; and
4 the storage controller including:
5 a first interface configured to receive block-level I/O requests;
6 a second interface configured for file-level I/O; and
7 a third interface configured for communication with another storage
8 controller,

9 wherein the controller presents a first logical volume by using first storage
10 area of the physical disk drives, and a second logical volume by using second storage area of
11 the physical disk drives,

12 wherein the controller receives a block-level write request to the first logical
13 volume at the first interface and then a file-level write request to a second logical volume at a
14 second interface,

15 wherein the controller performs a write operation to the first logical volume
16 and a write operation to the second logical volume, and

17 wherein the controller copies, through the third interface, write data
18 corresponding to the block-level write request to a third volume and then write data
19 corresponding to the file-level write request to a forth volume, the third volume and the forth
20 volume being presented by the another storage controller.

1 4. A method for operating a storage controller comprising:
2 receiving block-level I/O requests at a first interface;
3 accessing a physical data store to service the block-level I/O requests,
4 including accessing a first logical volume defined in the physical data store;
5 receiving file-level I/O requests at a second interface;
6 accessing the physical data store to service the file-level I/O requests,
7 including accessing a second logical volume defined in the physical data store; and
8 replicating the data stored in the first logical volume and in the second logical
9 volume, wherein the first logical volume and the second logical volume are defined as a
10 consistency group, wherein the step of replicating maintains time consistency of write
11 operations performed on the first logical volume and write operations performed on the
12 second logical volume.

1 5. The method of claim 4 wherein the first interface is configured for
2 communication with a SAN and the second interface is configured for communication with a
3 LAN.

1 6. A method for operating a storage controller connectable to a first
2 storage controller, comprising:

3 receiving a block-level write request to a first logical volume at a first
4 interface of the storage controller and then a file-level write request to a second logical
5 volume at a second interface of the storage controller, wherein the first logical volume and
6 the second logical volumes are presented by the storage controller;

7 performing a write operation to the first logical volume and a write operation
8 to the second logical volume; and

9 copying write data corresponding to the block-level write request to a third
10 volume and then write data corresponding to the file-level write request to a fourth volume,
11 wherein the third volume and the forth volume are presented by the first storage controller.

1 7. A storage system comprising:

2 a first interface configured to receive block-level I/O requests;

3 a second interface configured to receive file-level I/O requests; and

4 a data communication path suitable for data communication with physical
5 storage in order to exchange data with the physical storage in response to the block-level I/O
6 request and the file-level I/O requests,

7 wherein the block-level I/O requests are serviced by exchanging data with a
8 first storage area of the physical storage,

9 wherein write data to the first storage area is copied to a second storage area of
10 the physical storage,

11 wherein the file-level I/O requests are serviced by exchanging data with a third
12 storage area of the physical storage,

13 wherein write data to the third storage area is copied to a fourth storage area of
14 the physical storage,

15 wherein if a split request is received, then at least write data to the first storage
16 area that is received subsequent to receiving the split request is not copied to the second
17 storage area or write data to the third storage area that is received subsequent to receiving the
18 split request is not copied to the fourth storage area.

1 8. The storage system of claim 7 wherein if a re-sync request is received
2 then the first storage area is resynchronized with the second storage area and the third storage
3 area is resynchronized with the fourth storage area.

1 9. The storage system of claim 7 wherein if the first storage area and the
2 third storage area belong to a consistency group, then write data to the first storage area that
3 is received subsequent to receiving the split request is not copied to the second storage area
4 and write data to the third storage area that is received subsequent to receiving the split
5 request is not copied to the fourth storage area.

1 10. The storage system of claim 7 wherein the first interface is further
2 configured to communicate with a SAN.

1 11. The storage controller of claim 10 wherein the second interface is
2 further configured to communicate with a LAN.

1 12. A method for operating a storage controller comprising:
2 receiving one or more block-level I/O requests at a first interface;
3 receiving one or more file-level I/O requests at a second interface;
4 accessing a first logical volume defined in a physical data store to service the
5 block-level I/O requests;
6 mirroring block-level I/O write requests to a second logical volume defined in
7 the physical data store;
8 accessing a third logical volume defined in the physical data store to service
9 the file-level I/O requests; and
10 mirroring file-level I/O write requests to a fourth logical volume defined in the
11 physical data store.

1 13. The method of claim 12 further comprising:
2 receiving a split command;
3 if the split command is related to the first logical volume, then ceasing the
4 action of mirroring block-level I/O write requests; and
5 if the split command is related to the third logical volume, then ceasing the
6 action of mirroring file-level I/O write requests.

1 14. The method of claim 12 further comprising receiving a split command
2 and if the first logical volume and the third logical volume belong to a consistency group,
3 then ceasing the action of mirroring block-level I/O write requests and ceasing the action of
4 mirroring file-level I/O write requests.

1 15. A method for operating a storage controller comprising:
2 providing a first pair comprising a first logical volume and a second logical
3 volume, and a second pair comprising a third logical volume and a fourth logical volume;
4 receiving a block-level write request at a first interface;
5 receiving a file-level write request at a second interface;
6 performing a write operation to the first logical volume in accordance with the
7 block-level write request;
8 copying write data corresponding to the block-level write request to the
9 second volume;
10 performing a write operation to the third logical volume in accordance with
11 the file-level write request;
12 copying write data corresponding to the file-level write request to the fourth
13 volume; and
14 breaking at least one of the first pair and the second pair in accordance with a
15 split command.

1 16. The method of claim 15 further comprising:
2 if the split command is related to the first logical volume or to the second
3 logical volume, then ceasing the action of copying block-level I/O write requests; and
4 if the split command is related to the third logical volume or to the fourth
5 logical volume, then ceasing the action of copying file-level I/O write requests.

1 17. The method of claim 15 further comprising receiving a split command
2 and if the first logical volume and the third logical volume belong to a consistency group,
3 then ceasing the action of mirroring block-level I/O write requests and ceasing the action of
4 mirroring file-level I/O write requests.

1 18. A storage system comprising:
2 a controller;
3 a first physical storage connected to the controller;
4 a first interface configured to receive block-level I/O requests;
5 a second interface configured to receive file-level I/O requests; and
6 a third interface configured for communication with a storage subsystem, the
7 storage subsystem comprising second physical storage,
8 wherein the controller presents a first logical volume for the block-level I/O
9 requests by using a first portion of the first physical storage, and a second logical volume for
10 the file-level I/O requests by using a second portion of the first physical storage,
11 wherein the controller presents a third logical volume and a fourth logical
12 volume by using the second storage,
13 wherein the controller copies data in the first logical volume to the third
14 logical volume and copies data in the second logical volume to the fourth logical volume.

1 19. The storage system of claim 18 wherein the third logical volume is a
2 mirrored volume of the first logical volume, wherein the fourth logical volume is a mirrored
3 volume of the second logical volume, wherein if a split command is received, then write data
4 to the first logical volume is not copied to the third logical volume and write data to the
5 second logical volume is not copied to the fourth logical volume.

1 20. The storage system of claim 18 wherein the third logical volume is a
2 mirrored volume of the first logical volume, wherein the fourth logical volume is a mirrored
3 volume of the second logical volume, wherein if a split command is received, then either
4 write data to the first logical volume is not copied to the third logical volume or write data to
5 the second logical volume is not copied to the fourth logical volume.

1 21. A method for operating a storage controller of a first storage system
2 comprising:
3 providing a first logical volume and a second logical volume by using physical
4 storage included in the first storage system;
5 providing a third logical volume and a fourth logical volume by using physical
6 storage included in a second storage system, the second storage system being externally to
7 and in data communication with the first storage system;
8 receiving one or more block-level I/O requests to the first logical volume at a
9 first interface;
10 receiving one or more file-level I/O requests to the second logical volume at a
11 second interface;
12 performing write operations to the first logical volume in response to receiving
13 the block-level I/O requests;
14 performing write operations to the second logical volume in response to
15 receiving the file-level I/O requests;
16 copying write data corresponding to the block-level write requests to the third
17 volume; and
18 copying write data corresponding to the file-level write requests to the fourth
19 volume.

1 22. The method of claim 21 wherein the third logical volume is a mirrored
2 volume of the first logical volume, wherein the fourth logical volume is a mirrored volume of
3 the second logical volume, wherein if a split command is received, then write data to the first
4 logical volume is not copied to the third logical volume and write data to the second logical
5 volume is not copied to the fourth logical volume.

1 23. The method of claim 21 wherein the third logical volume is a mirrored
2 volume of the first logical volume, wherein the fourth logical volume is a mirrored volume of
3 the second logical volume, wherein if a split command is received, then either write data to
4 the first logical volume is not copied to the third logical volume or write data to the second
5 logical volume is not copied to the fourth logical volume.

1 24. A system comprising:

2 a first storage subsystem including physical storage, a first interface
3 configured to receive block-level I/O requests, and a second interface configured to receive
4 file-level I/O requests;
5 a second storage subsystem connected to the first storage subsystem, the
6 second storage subsystem managing a plurality of logical units; and
7 a backup server connected to the second storage subsystem,
8 wherein the first storage subsystem presents a first logical volume for the
9 block-level I/O requests by using the physical storage and a second logical volume for the
10 file-level I/O requests by using the physical storage,
11 wherein the first storage subsystem presents a third logical volume and a
12 fourth logical volume by using the logical units,
13 wherein the first storage subsystem copies data in the first logical volume to
14 the third logical volume and data in the second logical volume to the fourth logical volume,
15 wherein the backup server backs up data in the logical units.